

An Event Network for Exploring Open Information

Yanping Chen
ypench@gmail.com

No Institute Given

Abstract. In this paper, an event network is presented for exploring open information, where linguistic units about an event are organized for analysing. The process is divided into three steps: document event detection, event network construction and event network analysis. First, by implementing event detection or tracking, documents are retrospectively (or on-line) organized into document events. Secondly, for each of the document event, linguistic units are extracted and combined into event networks. Thirdly, various analytic methods are proposed for event network analysis. In our application methodologies are presented for exploring open information.

1 Introduction

Exploring information in open fields is a crucial and challenging task for human beings in the age characterized by flooding of information. Gigantic volumes of data spring up all over the world. This information is available for us by connecting into the Internet. Events (e.g. the Ebola virus, the Islamic state) outbreak far away from us may affect our ordinary life. We no longer turn a blind eye to world events, that we are getting involved. We not only consume information, but also produce it. Social media (e.g., Twitter) makes us easily and freely express opinions. Millions of comments spread over the Internet, affecting the trend of public opinions. Accessing this information is beneficial for our decision-makings. Commonly, Information Retrieval (IR) and Information Extraction (IE) provide helpful solutions for us to handle the surge of information.

When exploring information, event oriented techniques provide effective approaches to understand who, where, when and what happened. Output of event detection in IR is series (or clusters) of documents, referred as *document events* in this paper. The problem with IR systems is that, they only retrieve subject-documents, but no content of documents is identified. Users are required to skim through returned documents. Events of IE are templates with slots to be filled, referred to as *template events*. Template event recognition extracts structured data from semi-structured or unstructured data. It is expected that the result is used to populate a knowledge base directly. The main problem for template event recognition is that it suffers from poor performance, especially in an open field, where heterogeneous resources, noise and fragmental data are processed.

In open fields, there are knowledge bases automatically constructed to support information exploring. Many of them are constructed from semi-structured database (e.g. Wikipedia, WordNet) or under human supervision (collaboratively). Therefore, a better consistency is expected. They are widely used as external knowledge for semi-supervised methods. Because semi-structured data or human labor are required, when a

news is broken in open fields automatically handling this information in real time is difficult. Furthermore, many systems automatically organize extracted linguistic units into a graph-based representation. These systems generally lead to a complex network with thousands of nodes or edges. Rare analysis was conducted to show underlying structures of events. Generally, output generated by IE systems is error-prone, redundant and incompatible information makes it contradictory. In this paper, analyses of event network are emphasized. Contributions of this paper include,

1. An event network framework is proposed, which provides an event oriented information exploring and supports various analyses.
2. Two novel analytic methods (PLT analysis and action analysis) are discussed for event oriented information exploring.

The rest of this paper is organized as follows. Section 2 introduces related work. Motivations and definitions about the event network are presented in Section 3. Section 4 discusses our method to construct event networks. Section 5 demonstrates applications of event network. Section 6 gives conclusions.

2 Related Work

Organizing linguistic units as networks or graphs shows an increased interest in NLP. These representations can be roughly divided into three paradigms: logic based semantic network (e.g. Ontology), scalable knowledge database (e.g. Freebase) and semantic network constructed from open information extraction.

Logic based semantic network refers to networks constructed with human labor. They mainly focus on a closed domain. For example, *conceptual graph* represents logic as a graph representation (Sowa, 1984). These networks support logic operators, can map questions and assertions from natural language to a relational database. Logic based semantic networks are constructed by domain experts and used as domain-specific ontologies (e.g., WordNet, Cyc, etc.). Commonly, they support inferences developed in knowledge representation. In this paradigm, conflicting and contradictory are not allowed. Therefore, in open fields it is difficult to apply.

In the second paradigm, logic based semantic network is extended into open fields. Large knowledge databases such as Yago (Suchanek et al., 2007), Freebase (Bollacker et al., 2008) are constructed, representing large knowledge in formalized forms. These representations merge diverse and heterogeneous data with high scalability, providing a unified framework for organizing information. Instead of rigid definitions, these networks have no canonical view of data. They use a loose representation to support scalability and extensibility. Many of them are constructed by merging ontologies (e.g. WordNet, OpenCyc) or extracted from semi-structured database (e.g. Wikipedia). To construct logic based semantic network, commonly, direct or indirect human interventions are involved (e.g. collaborative methods or searching logs).

Instead of aiming at semi-structured data, the third paradigm explores information in an open and dynamic field (mainly focusing on unstructured data). In this paradigm, weak supervision (Mintz et al., 2009; Xu and Zhao, 2014) and bootstrapping methods (Agichtein and Gravano, 2000; Kozareva and Hovy, 2010; McIntosh et al., 2011;

Weld et al., 2009) are employed. Scalable knowledge databases are used to guide the process, e.g., TEXTRUNNER (Banko et al., 2007), KNOWITALL (Etzioni et al., 2005, 2011), WOE (Hoffmann et al., 2010) and StatSnowBall (Zhu et al., 2009). Commonly, in these systems, nodes are named entities and edges are relations between them. All extracted results are combined into a large network. It often generate a network containing more than thousands of nodes and edges.

The notion of event is widely used for exploring information. Piskorski et al. (2011) presents an on-line news event extraction system. Each event is defined as a frame with slots filled by information extracted from clustered documents, where the pattern matching method is used. Ramakrishnan et al. (2014) proposes an EMBERS system encoding events as frames. It is used to forecast “civil unrest” events in open fields. TwiCal extracts open-domain events from Twitter (Ritter et al., 2012), where events are identified by named entities. In the same data set, ET represents events as clusters of keywords (Parikh and Karlapalem, 2013). Kuzey et al. (2014) uses events itself as nodes of network. They cluster documents into a hierarchical representation, where nodes are events and edges link the same event in chronological order. Angel et al. (2012) constructs an entity network from social media by the *streaming edge weight* method. They mine dense subgraphs of network for identifying realtime stories. Das Sarma et al. (2011) provides an event discovery method based on entity dynamic relation graphs, which are constructed by co-occurrences of entities constrained in documents.

3 Motivation and Definition

Methodologies to organize linguistic units into networks or graphs show an increased interest in NLP. They provide novel solutions for many NLP tasks and support human oriented information exploring. Representing linguistic units as a graph enables topological analyses developed in fields such as: social network and complex network.

In open fields, these representations are mainly constructed by techniques developed under information extraction or text understanding. Information extraction aims at extracting linguistic units with concrete concepts or functions. It is seen as a trade-off between information retrieval and text understanding, where text understanding tries to capture all information in a document. Text understanding may lead to worsen performance caused by applied techniques (Hobbs and Riloff, 2010). On the other hand, information extraction extracts targeted units and ignores uninterested.

Due to extracting challenges in open fields, instead of extracting information in a monolithic process, we divide the task into three steps: document event detection, event network construction and event network analysis. In the first step, by implementing document event detection and tracking, documents are organized into document events. Most of irrelevant or uninterested information is filtered. Then, in the second step, IE techniques are employed to extract linguistic units and organized into event networks. Techniques with higher performance are highlighted. For example, instead of extracting named entities as nodes, entity mentions are used. Where coreference resolution are required to group entity mentions into named entities, which is error-prone. In the last step, because topological information is available, structural information between lin-

guistic units can be used to modify the network quality. Then, network or graph based analytic methods can be used, and it is convenient for visualization.

In this domain, many systems combine extracted results into a complex network, which not doing much help for analysing information. Furthermore, redundant and incompatible information makes it contradictory and misunderstanding. In our application, we emphasize methodologies conducted for event network analysis. Advantages of the event network include: first, after document event detection, information extraction in each document event can be independently implemented. Therefore, the effect of noise and heterogeneous data on information extraction can be reduced. Secondly, crossing document information enables discovery of potential relations between documents. Thirdly, event network provides a structured data representation for exploring open information, topological methods, e.g., social network or complex network, can be introduced for event network analysis.

For convenience to discuss event network analysis, we define *nodes* and *edges* of event network as *frames* with *slots*. These information is also expected to support human oriented information exploring.

Let $\mathcal{D}=\{d_1, \dots, d_L\}$ be a document set, d_i denotes a document. A document event \mathcal{E}_k is a subset of \mathcal{D} . For all $d_i, d_j \in \mathcal{E}_k$, similarity function $Similarity(d_i, d_j)$ satisfy a predefined condition (e.g. a threshold). All document events in \mathcal{D} are denoted as $\mathcal{E}=\{\mathcal{E}_1, \dots, \mathcal{E}_K\}$. The constraint that \mathcal{E} is a partition of \mathcal{D} is not necessary, because some documents in \mathcal{D} can be filtered, or *fuzzy partitioning techniques* can be used, which enable a document belonging to more than one document event.

An event network on document event \mathcal{E}_k is represented as a graph $\mathcal{N}_k=\{V_k, E_k\}$, where $V_k=\{v_{k1}, \dots, v_{kN}\}$ and $E_k=\{e_{k1}, \dots, e_{kM}\}$ are vertex set and edge set. Both vertices and edges are frames defined as follows.

$$vertex := \{key, name, type, weight, info\}$$

$$edge := \{type, v-1, v-2, weight, info\}$$

where *vertex* frame defines nodes of event network. Slot “*name*” refers to entity mentions occurred in a document event. Each vertex is identified by an integer value “*key*”. Slot “*type*” represents categories of vertices (e.g. *Person*, *Organization* and *Location*). Slot “*weight*” is the likelihood of “*name*” to be “*type*”. Traditionally, this value is given by a classifier when extracting this frame. Depending on real applications, “*weight*” can be used to filter an event network. An *edge* frame denotes a relation between two vertexes. Slots “*v-1*” and “*v-2*” are *keys* of vertexes in an edge, used to identify vertexes linked by edges. Edge types are referred by “*type*” (e.g. Part-whole, Personal-Social). In both frames, slot “*info*” contains information about the frames, where entity mentions or entity relations occurred, e.g., sentences, documents or timestamps. These information support event network analyses (e.g. coreference resolution, statistical relational learning or manually exploring). If they are empty, these values are *null*.

4 Implementation

This section discusses our method to construct event networks, which are used to show methodologies discussed in Section 5.

4.1 Data Sets

We use the ACE 2005 Chinese corpus. It contains 633 documents annotated with 15,264 entities and 33,932 entity mentions¹. 7 entity types (e.g. *person*, *organization*, etc.) and 44 entity subtypes are defined. The corpus also annotated with 6 major relation types and 18 relation subtypes. Each relation instance has two named entities as arguments. There are 9,244 relation mentions are collected as positive instances.

The ACE 2005 Chinese corpus is used to train named entity and relation classifiers. In order to show our method in an open field, we also use the Chinese Gigaword Fifth Edition corpus. The Peoples Daily source is used, which contain 145,001 newswire texts covering the period from November 2006 through December 2010.

4.2 Events Detection

The purpose of document event detection is to cluster documents into events. We use LDA toolkit provided by Phan and Nguyen (2007) to implement this task. In LDA model, a corpus is first represented as a matrix, where each column refers to a document vector, and each row represents distribution of a term in documents. Then LDA maps documents from a term space into a topic space. Topics are hidden variables.

Because we focus on newswire texts, where short texts are commonly used. We use Omni-word feature proposed by Chen et al. (2014), which takes every potential word as terms of documents. It is a subset of n-Gram feature. In the pretreatment process, we remove high and low frequency words² in an employed lexicon. Words with frequencies lower than 10 are also removed. To train an LDA model, hyper-parameters are required. The topic number is set as 25. Other parameters use default settings.

The toolkit generates several outputs. The word-topic distributions are more favourable to us, which give distributions of terms in a topic space. We use topics as centroids of document clusters in a term space. When clustering documents, a documents belonging to an event is judged by the nearest *Euler Distance* of the document and centroids. The top 100 most likely words per each topic are used to represent an event.

It is recognized that documents discussing the same event tend to be temporal proximity, and a time gap between bursts of similar documents may indicate different events (Yang et al., 1999). Therefore, timestamps are used to partition the newswire texts. In our experiment, the time step is set as 5 months. Then, the Chinese Gigaword corpus is divided into 10 parts. Each part contains 5 months newswire texts. Because hierarchical representation can give a multi-granularity review when exploring open information and reduce the travel cost. In each time step, instead of using retrospective methods to give a flat partition of documents, we organize them into a hierarchical representation. Documents of each time step are clustered into 25 events by the LDA toolkit. Each event is further clustered into sub-events by the same approach. If an event contains documents less than ten, the process to find its sub-events is skipped. Therefore, in each time step, 25 events and at most 25×25 sub-events are detected.

¹ An *entity mention* is a reference to an entity.

² The ratio is 5% for each.

4.3 Named Entity Recognizing

In this step, it is free to use any named entity recognition methods. In our application, we use a Boundary Assembling (BA) method to implement the named entity recognition task. The notion of BA method is that, instead of recognizing entity mentions in a unitary style, it first detects boundaries of entity mention, then assembles detected boundaries into entity mention candidates. Each candidate is further assessed by a classifier.

In our work, we recognize three types of named entity: “PER” (Person), “LOC” (Location) and “ORG” (Organization). In order to filter noise, recognized named entities with Chinese characters less than two and more than six are discarded.

4.4 Relation Recognizing

To recognize relations between named entities, we adopt the method proposed in Chen et al. (2014), where an Omni-word feature and a soft constraint method is proposed for Chinese relation extraction. The Omni-word feature uses every potential word in a relation mention as lexical features. Then for each employed atomic feature, an appropriate constraint condition is selected to combine them with additional information to maximize the classification determination.

With our employed three entity types, five relation types annotated in the ACE corpus are recognized: “PER-SOC”, “GEN-AFF”, “ORG-AFF”, “PART-WHOLE” and “PHYS”. Sentences with more than ten entities are ignored, because extracting relations in a long sentence is error-prone.

4.5 Merging and Visualizing

As approaches discussed above, the result about recognized document events, named entities and relations are listed in Table 1.

Table 1: Information of Results

Step	Doc.	Event	Entity	Relation	Step	Doc.	Event	Entity	Relation
0	14,814	642	1,123,506	309,847	5	19,706	643	1,389,668	290,594
1	11,734	642	897,489	258,817	6	12,014	645	927,977	265,439
2	17,678	644	1,104,517	257,291	7	9,100	641	705,805	254,133
3	18,213	646	1,305,191	277,721	8	9,326	635	726,123	230,535
4	23,433	644	1,704,294	289,166	9	8,983	640	674,457	243,210

In the follows, in order to conduct event network analysis, we use the igraph toolkit provide by Csardi and Nepusz (2006), representing extracted entity mentions and relations as a graph. A network analysis toolkit (Pajek) provided by Batagelj and Mrvar (1998) is used for visualization.

5 Application

In our work, we emphasize analyses of event networks. After event networks were constructed, techniques such as social network, complex network can be employed to analyse event networks. For example, setting a person name as a central entity, we can navigate entities around it. Filtering irrelevant information, we can show character relationships in an event. Using the “PART-WHOLE” relations, multi-granularity visualization can be supported. Because event networks have a graph representation, topological information is available. Therefore, various approaches (e.g., statistical relational learning) can be used to improve the network quality. Furthermore, event networks support human oriented information exploring. When human exploring open information, manual interventions can be used to modify the quality of event networks.

In this section, we choose a document event in time step 0 as an example, which contain 1,041 documents. There are 42,436 named entities and 6,272 relations occurred. The most likely words in it are “袭击,北约,发言人,冲突,防御, etc.” (Assault, NATO, Spokesman, Conflict, Defence, etc.). It indicates that the concern of this event is military affairs. Extracted named entities and relations are organized in Figure 1, where there are 252 nodes and 571 edges are merged. Nodes in *Red*, *Yellow* and *Blue* colors represent *Person*, *Organization* and *Location* respectively. Each edge is labelled by the relation type.





Fig. 1: Event Network

To explore open information, many systems dynamically organize linguistic units into a complex network. Heterogeneous resources and unreliable information make the network chaotic and misunderstanding. As Figure 1 showing, a complex network makes it difficult to understand. In the following, based on the event network, we give four methodologies to explore open information: Information Filtering, PLT Analysis, Action Analysis and Social Network Analysis.

5.1 Information Filtering

The simplest way to analyse event network is to filter information that is irrelevant or uninterested. In Figure 2(a), only person names and “PER-SOC” relations are remained to show character relationships in an event network.



This example can be formalized as: Let \mathcal{N} be an event network. The filtered event network $\mathcal{N}' = \{V', E'\}$ is a subgraph of \mathcal{N} , such that $\mathcal{N}' \subset \mathcal{N}$. And \mathcal{N}' satisfies $\forall v \in V' (v.type = \text{PER}) \wedge \forall e \in E' (e.type = \text{PER-SOC})$.

Using information contained in *vertex* frames and *edge* frames, information filtering can provide effective approaches for exploring open information. For example, in *vertex* frames and *edge* frames, we may require that the value in *weight* slots is greater than a predefined threshold. Utilizing information in *info* slots, we can collect named entities occurred in specified periods or areas. For a central figure, we can see directly connected named entities and relations between them.

5.2 PLT Analysis

Person-Location-Time (PLT) analysis tries to find relations between persons and locations in a period of time. It can be used to track a person, find trajectories of targeted entities. *Person*, *Location* are extracted by named entity recognition methods. While the *Time* is different. Two kinds of *Time* are distinguished in a document: *implicit temporal information* and *explicit temporal information*. Implicit temporal information is part of the document’s content indicating the creation, development, termination of an event. It is also seen as a named entity type in some researches. Extracting this information needs information extraction or text understanding techniques. In many applications, it is ignored. Generally, in open fields, all documents have explicit temporal information, which includes the creation, modification and transmission timestamps of documents. They are meta-data spread with documents. In our paper, because we focus on newswire texts, where the explicit temporal information of documents is released together. Therefore, we use the explicit temporal information for PLT analysis.

This process can be formalized by introducing an attribute *time* in the *info* slot. Let \mathcal{N} be an event network, *times* represent timestamps and *person* is a person name. $\mathcal{N}' = \{V', E'\}$ is the result of PLT analysis based on \mathcal{N} , where $\forall e \in E' (e.type = \text{PHYS} \wedge (e.v-1 = \text{person} \vee e.v-2 = \text{person}))$. In other words, all relation type in E' is “PHYS”, and take the same entity mention *person* as an argument. Replacing all *person* by corresponding *times*, we get a graph with nodes referred to timestamps and locations. An example is shown in Figure 2(b).

In this example, we track Mao Zedong (“毛泽东”)³ in the whole Gigaword corpus, collect all recognized “PHYS” relation instances which have Mao Zedong as an argument. In the result, there are 142 “PHYS” relation mentions, which take Mao Zedong (or Chairman Mao) as arguments. Then we replace Mao Zedong (or Chairman Mao) by the explicit temporal information of newswire texts. In Figure 2(b)⁴, nodes in green color are timestamps, and blue nodes are locations. Each green node means that Mao Zedong occurred with the connected locations at that time.

5.3 Action Analysis

Recognizing an “event” under the ACE definition is difficult, where event triggers, participant roles, properties and attributes should be identified (Doddington et al., 2004). It received an ACE value score only 30% in Ahn (2006). In an open field, it will come to worse performance. In many researches, co-occurrence information (e.g. co-citation, co-word, co-link, etc.) between terms is used to explore and understand structures in the underlying document sets, e.g., Leydesdorff and Vaughan (2006). In our application, instead of the definition in ACE, we present the action analysis.

In action analysis, we focus on detecting whether or not a special action is mentioned in a sentence. Therefore, we conduct the “sentence classification” task, classifying each sentence by a classifier trained on the ACE annotated event mentions. In our application, we monitor the “Conflict” ACE event type, which has two subtype: Attack and Demonstrate (Doddington et al., 2004). The ACE corpus, which annotates 596 “Conflict” events, is employed for training and testing. We implement the 5-fold cross validation, and the P/R/F (Precision/Recall/F-score) measurement. F-score is computed by $(2 \times P \times R) / (P + R)$. In order to perform a two-class classification, we generate negative instances by segmenting the corpus into sentences, discarding annotated ACE event mentions, and filtering sentences without event triggers of “Conflict” ACE events. Then, 1,589 sentences are collected as negative instances. We only use Omniwords features in sentences for classification. The performance is shown in Row 1 of Table 2, where Only the performance about “Conflict” is listed.

Table 2: Performance of Action Analysis

No.	Precision	Recall	F-core
1	82.19	82.88	82.53
2	97.65	41.94	58.68

In an open field with massive data, the precision is more emphasized. Therefore, we label an instance as a “Conflict” action only when the employed classifier (maximum entropy) output a predicted value equals 1⁵. The performance is shown in Row 2 of

³ The leader of the Communist Party of Chinese.

⁴ Because the original graph is more complex (55 nodes and 142 edges), in this place, only part of it is given.

⁵ The default value is 0.5 in two-class classification.

Bibliography

- Eugene Agichtein and Luis Gravano. Snowball: Extracting relations from large plain-text collections. In *Proceedings of DL '00*, pages 85–94. ACM, 2000.
- David Ahn. The stages of event extraction. In *Proceedings of the Workshop on Annotating and Reasoning about Time and Events*, pages 1–8. ACL, 2006.
- Albert Angel, Nikos Sarkas, Nick Koudas, and Divesh Srivastava. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. *PVLDB*, 5(6):574–585, 2012.
- Michele Banko, Michael J Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. Open information extraction for the web. In *Proceedings of IJCAI '07*, volume 7, pages 2670–2676, 2007.
- Vladimir Batagelj and Andrej Mrvar. Pajek-program for large network analysis. *Connections*, 21(2):47–57, 1998.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of SIGMOD '08*, pages 1247–1250. ACM, 2008.
- Yanping Chen, Qinghua Zheng, and Wei Zhang. Omni-word Feature and Soft Constraint for Chinese Relation Extraction. In *Proceedings of ACL'14*, pages 572–581. ACL, 2014.
- Gabor Csardi and Tamas Nepusz. The igraph software package for complex network research. *IJ COMP SYS*, 1695(5):1–9, 2006.
- Anish Das Sarma, Alpa Jain, and Cong Yu. Dynamic relationship and event discovery. In *Proceedings of WSDM '11*, pages 207–216. ACM, 2011.
- G. Doddington, A. Mitchell, M. Przybocki, L. Ramshaw, S. Strassel, and R. Weischedel. The automatic content extraction (ACE) program—tasks, data, and evaluation. In *Proceedings of LREC '04*, volume 4, pages 837–840. Citeseer, 2004.
- O. Etzioni, M. Cafarella, D. Downey, A.M. Popescu, T. Shaked, S. Soderland, D.S. Weld, and A. Yates. Unsupervised named-entity extraction from the web: An experimental study. *AI*, 165(1):91–134, 2005.
- Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, and Mausam Mausam. Open Information Extraction: The Second Generation. In *Proceedings of IJCAI '11*, volume 11, pages 3–10, 2011.
- Jerry R Hobbs and Ellen Riloff. Information extraction. *Handbook of natural language processing*, 2, 2010.
- R. Hoffmann, C. Zhang, and D.S. Weld. Learning 5000 relational extractors. In *Proceedings of ACL '10*, volume 10, pages 286–295. ACL, 2010.
- Zornitsa Kozareva and Eduard Hovy. Learning arguments and supertypes of semantic relations using recursive patterns. In *Proceedings of ACL'10*, pages 1482–1491, 2010.
- Erdal Kuzey, Jilles Vreeken, and Gerhard Weikum. A Fresh Look on Knowledge Bases: Distilling Named Events from News. In *Proceedings of CIKM '14*, pages 1689–1698. ACM, 2014.

- Loet Leydesdorff and Liwen Vaughan. Co-occurrence matrices and their applications in information science: Extending ACA to the Web environment. *JASIST*, 57(12): 1616–1628, 2006.
- Tara McIntosh, Lars Yencken, James R Curran, and Timothy Baldwin. Relation Guided Bootstrapping of Semantic Lexicons. In *Proceedings of ACL '11*, pages 266–270, 2011.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. Distant supervision for relation extraction without labeled data. In *Proceedings of ACL '09*, pages 1003–1011. ACL, 2009.
- Ruchi Parikh and Kamalakar Karlapalem. Et: events from tweets. In *Proceedings WWW '13*, pages 613–620. IW3C2, 2013.
- Xuan-Hieu Phan and Cam-Tu Nguyen. Gibbslda++: Ac/c++ implementation of latent dirichlet allocation, 2007.
- Jakub Piskorski, Hristo Tanev, Martin Atkinson, Eric Van Der Goot, and Vanni Zavarella. Online news event extraction for global crisis surveillance. In *TCCI*, pages 182–212. Springer, 2011.
- Naren Ramakrishnan, Patrick Butler, Sathappan Muthiah, Nathan Self, et al. 'Beating the news' with EMBERS: Forecasting Civil Unrest using Open Source Indicators. In *Proceedings SIGKDD '14*, 2014.
- Alan Ritter, Oren Etzioni, Sam Clark, et al. Open domain event extraction from twitter. In *Proceedings of SIGKDD '12*, pages 1104–1112. ACM, 2012.
- John F Sowa. *Conceptual structures: information processing in mind and machine*. 1984.
- Fabian M Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: a core of semantic knowledge. In *Proceedings of WWW '07*, pages 697–706. ACM, 2007.
- Daniel S Weld, Raphael Hoffmann, and Fei Wu. Using wikipedia to bootstrap open information extraction. *ACM SIGMOD Record*, 37(4):62–68, 2009.
- Yang Liu Kang Liu Liheng Xu and Jun Zhao. Exploring Fine-grained Entity Type Constraints for Distantly Supervised Relation Extraction. 2014.
- Yiming Yang, Jaime G Carbonell, Ralf D Brown, Thomas Pierce, Brian T Archibald, and Xin Liu. Learning approaches for detecting and tracking news events. *IEEE INTELL SYST*, 14(4):32–43, 1999.
- Jun Zhu, Zaiqing Nie, Xiaojiang Liu, Bo Zhang, and Ji-Rong Wen. StatSnowball: a statistical approach to extracting entity relationships. In *Proceedings of WWW '09*, pages 101–110. ACM, 2009.

Document Event Viewer			X
ID	Key Words	Sul	^
9	选举;立法;内地;特区;人权;听取;法制;妇女;提案;修改;少数;同胞;议案;区政府;座谈;座谈会;公民;班子;...	T	
10	食品;环保;节约;施工;污染;仓库;加工;供应;建筑;销售;伙食;安装;改造;蔬菜;燃料;发电;电站;垃圾;广西;...	T	
11	你们;老人;父亲;儿子;心里;母亲;下来;父母;家人;家里;家乡;妻子;一下;晚上;怎么;故事;回来;感动;她们;...	T	
12	民法;司法;人民法院;民工;被告;职务;判处;审判;检察院;处罚;徒刑;高人;刑事;兴奋;判决;赔偿;纠纷;兴...	T	
13	试飞;海上;战机;坦克;空中;雷达;潜艇;直升机;舰艇;试验;地面;性能;飞行员;侦察;战舰;攻击;海域;轨道;...	T	
14	电磁;模式;指挥员;战时;机械;评估;模拟;检验;新装;干扰;物资;兵种;一体化;机械化;课目;数字;变革;精...	T	
15	艺术;文艺;表演;音乐;节目;节日;演出;作品;创作;感受;作者;军旅;精彩;新年;网友;新春;热烈;照片;歌曲;...	T	
16	警方;嫌疑;非法;嫌疑人;保密;民警;治安;犯罪嫌疑人;联网;浙江;抓获;公安局;手机;毒品;互联网;上网;...	T	
17	煤矿;天津;明星;篮板;救援;受伤;联盟;西部;恢复;得分;赛区;热火;最近;太阳;最佳;板球;篮板球;工人;助...	T	
18	志愿;市民;少年;参观;网站;残疾;展览;奥委会;组委会;健身;展示;场馆;主办;老年;理事;残疾人;青少年;...	T	
19	支部;总队;安置;老干部;特约;党支部;驻地;通讯;家属;军地;下发;退役;做法;通讯员;选取;出台;年初;工...	T	
20	生态;小学;面积;援建;气象;贫困;蒙古;四川;村民;山东;天然;支援;自治区;贵州;内蒙;内蒙古;平方;森林;...	T	
21	往往;容易;其实;批评;那么;比如;思维;如此;那些;带兵;无论;只是;人生;自我;态度;在于;缺乏;追求;习惯;...	T	
22	峰会;交往;对话;东盟;共识;密切;个中;达成;平等;增进;越南;高层;朋友;建交;签署;边关;繁荣;评价;务实;...	T	
23	比分;主教;主教练;主场;球迷;世界杯;女排;意大利;击败;网球;女足;国家队;进攻;开赛;失误;欧洲;男篮;...	T	
24	袭击;北约;发言人;冲突;防御;官员;南部;恐怖;威胁;军方;声明;北部;局势;欧洲;日报;地雷;欧盟;反对;爆...	T	
<div> <div>Father</div> <div>Children</div> <div>Hierarchy: 1/3</div> <div>Doc Num: 1041</div> <div>Events Num: 25</div> <div>Load Events</div> <div>Activate Event Analysis</div> </div>			
<div> <div>Timestamp: <<</div> <div>Timestamp-200611</div> <div>>></div> </div>			Selected Event: Event-0-024